REMARKS

First note

Applicant has cancelled claims 7-10, 12-15, 16-19, and 21-24 in this response after final. Applicant respectfully requests that the Examiner enter this amendment, because it places the present patent application in better condition for appeal. No additional search or consideration is required by the Examiner to enter this amendment, since the only amendment being proffered herein is the cancellation of claims.

Because claims 7-10, 12-15, 16-19, and 21-24 have been cancelled, the rejections and objections relating thereto are effectively moot, and Applicant only discusses below the rejections and objections in the final office action as they pertain to the remaining pending claims, claims 1, 4-6, 25, and 28-30.

Second note

Applicant is prepared to appeal the present patent application. However, Applicant would like to give the Examiner the opportunity to suggest claim amendments that would render the application allowable. Applicant believes that Applicant and the Examiner are "close" in reaching a "meeting of the minds" as to the patentability of the claims, and that the primary issues of disagreement are more in the way of form instead of substance. Therefore, if the Examiner believes likewise, he is respectfully encouraged to contact Applicant's representative, Mike Dryja, at the phone number listed below, with any suggested claim amendments.

Claim rejections under 35 USC 112

Claims 1 and 25 have been rejected under 35 USC 112, second paragraph, as being indefinite, because the Examiner has stated that these claims are directed to both an apparatus and method steps of using the apparatus, citing IPXL Holdings v. Amazon, 430 F.2d 1377, 1384 (Fed. Cir. 2005) in particular. The Examiner has noted that the offending phrase in question in

this respect is the phrase "even if the second node has a cache, memory, and at least one processor" found in both of these claims. Applicant respectfully disagrees.

First, claims 1 and 25 are not directed to an apparatus and method steps of using an apparatus. Rather, claims 1 and 25 are directed to a method having method steps. It is noted that in the IPXL Holdings case, the claim at issue recited the following:

> 25. The system of claim 2 [including an input means] wherein the predicted transaction information comprises both a transaction type and transaction parameters associated with that transaction type, and the user uses the input means to either change the predicted transaction information or accept the displayed transaction type and transaction parameters.

Thus, the Federal Circuit concluded in this case that the claim was directed to an apparatus – a "system" – as well as to method steps of using the apparatus – the "user *uses* the input means [of the system]." There is no correspondence as to this holding and the present patent application. Claims 1 and 25 of the present patent application are directed to a method having method steps, and not to an apparatus and method steps of using the apparatus. That is, claims 1 and 25 are not directed to an apparatus at all, but rather are directed to "a method comprising," as opposed to "an apparatus comprising" or "a system comprising," as in IPXL Holdings.

Second, claims 1 and 25 do not actually claim that a "second node [that] has a cache, memory, and at least one processor," and thus cannot be considered as being directed to an apparatus, or claiming an apparatus, in any way. The complete claim limitation is that "if the first node to which the given I/O device is connected has a cache, memory, and at least one processor, then assigning the given interrupt for the given I/O device to the first node, even if the second node has a cache, memory, and at least one processor." This is a method step, pure and simple, without any claiming of the second node and its constituent components.

That is to say, Applicant is claiming that *even if* the second node has a cache, memory, and at least one processor, the method *still* assigns the given interrupt to the first node, which is a method step. Claims 1 and 25 do not claim a second node that has a cache, memory, and at least one processor, but rather is claim a method *that is performed in relation to such a second node*.

Claiming a method and method steps that are performed in relation to an apparatus does not elevate claims 1 and 25 to claiming the apparatus itself.

Finally, third, Applicant requests that the Examiner consider the following representative claim language "playing baseball even if the clouds release rain." The claim recites a method step – playing baseball – which is subject to the condition that baseball will be played "even if the clouds release rain." The claim is not directed to clouds, and the claim does not actually claim clouds. Likewise, assigning a given interrupt in the claimed invention is achieved "even if the second node has a cache, memory, and at least one processor." Claims 1 and 25 are not directed to the second node, and the claim does not actually claim the second node.

For all of these reasons, Applicant respectfully submits that claims 1 and 25 are not indefinite under 35 USC 112, second paragraph, because these claims 1 and 25 recite and are purely directed to a method and its constituent method steps, and are not directed to nor claim an apparatus in anyway.

Claim rejections under 35 USC 101

The Examiner has stated that "the claimed invention" is directed to non-statutory subject matter, because "the claim" is directed to neither a process nor a machine, but rather embraces two different statutory classes of the invention. In the first instance, Applicant respectfully requests that the Examiner identify in an advisory action which claims he is referencing in this rejection. For the purposes of this rejection, Applicant presumes that the Examiner is discussing claims 1 and 25 (although this is unclear, since the Examiner refers to "the claim" singular, and not to "the claims" plural). In the second instance, insofar as claims 1 and 25 are not directed to both a process and a machine, as has been discussed above in relation to 35 USC 112, second paragraph, above, these claims satisfy 35 USC 101.

Stated another way, it appears that the rejections under 35 USC 112, second paragraph, and the rejections under 35 USC 101 are inextricably linked. Insofar as the claims satisfy or fail

to satisfy 35 USC 112, second paragraph, they likewise satisfy or fail to satisfy 35 USC 101. Applicant cannot envision any situation in which claims 1 and 25 fail to satisfy 35 USC 112, second paragraph, per the Examiner's reasoning discussed above while nevertheless satisfying 35 USC 101, or vice-versa. Therefore, Applicant respectfully submits that insofar as claims 1 and 25 satisfy 35 USC 112, second paragraph, as noted above, these claims also are statutory under 35 USC 101.

Claim objections

Claims 1 and 25 have been objected to because the Examiner has stated that they are "of improper dependent form for failing to further limit the subject matter of a previous claim." The Examiner has stated that "Applicant is required to cancel the claims, or amend the claims to place the claims in proper dependent form, or rewrite the claims in independent form." Applicant notes, however, that claims 1 and 25 are *already* in independent form – they are both independent claims. Therefore, these objections do not make sense. That is, claims 1 and 25 cannot be "of improper dependent form for failing to further limit the subject matter of a previous claim," because they are not in dependent form, and because they are in independent form, by definition they inherently cannot further limit the subject matter of a previous claim. Likewise, the Examiner has stated that one way to overcome this rejection is to rewrite claims 1 and 25 in independent form – but claims 1 and 25 are already in independent form. Applicant is very confused, and asks that the Examiner either withdraw these objections, or issue a new final office action to which Applicant can properly respond. (It is noted that clarifying this objection in an advisory action would be improper, since Applicant would not then have had an opportunity to seasonably respond to the Examiner, because a response cannot be filed after an advisory action.)

The Examiner has stated that the substance of these objections is that claims 1 and 25 recite "even if the second node has a cache, memory, and at least one processor." The Examiner has stated that this "limitation is essentially the same as 'whether or not' and thus fails to further

limit the claim, since the step is carried out regardless." In this respect, Applicant very much disagrees with the Examiner. To make this clear, Applicant will use some shorthand. The phrase "if the first node to which the given I/O device is connected has a cache, memory, and at least one processor" is abbreviated as "if A"; the phrase "if the second node has a cache, memory, and at least one processor" is abbreviated "if B"; and the phrase "assigning the given interrupt for the given I/O device to the first node" is abbreviated "perform C." Therefore, claims 1 and 25 *prior to amendment in the previous office action response*, recite "if A then perform C," whereas after amendment in the previous office action response, these claims recite "if A then perform C even if B."

Now, consider the following hypothetical prior art that teaches "if A then perform C" and "if A and B then perform D". This prior art reads on claims 1 and 25 prior to amendment, because claims 1 and 25 prior to amendment recite "if A then perform C" and the hypothetical prior art likewise teaches "if A then perform C." The Examiner would properly reject these claims over the hypothetical prior art in this instance. However, this prior art does *not* read on claims 1 and 25 after amendment, because claims 1 and 25 after amendment recite "if A then perform C even if B," and if both A and B occur/exist, the prior art teaches that D is performed (not C as in claims 1 and 25), since the prior art teaches "if A and B then perform D."

Therefore, it cannot be said that the limitation "even if B" fails to further limit the claim. In fact, this limitation does further limit claims 1 and 25. In the examples noted in the previous paragraph, the hypothetical prior art reads on claims 1 and 25 without this limitation, but does not read on these claims with this limitation. The only conclusion one can logically draw, then, is that the limitation "even if B" does indeed further limit the claims, because the scope of the claims changes due to its inclusion – as evidenced by the fact that the hypothetical prior art reads on the claims without this limitation but does not read on the claims with this limitation.

Applicant notes in this respect that "if A then perform C even if B" is tantamount to reciting "if A and B then perform C" and "if A and not B then perform C." These two statements

are not the same as just stating that "if A then perform C," however. Consider again the hypothetical prior art that teaches "if A then perform C" and "if A and B then perform D." Prior to amendment, claims 1 and 25 recited "if A then perform C," and thus the hypothetical prior art reads on them, since the hypothetical prior art explicitly teaches "if A then perform C." By comparison, after amendment, claims 1 and 25 equivalently recite "if A and B then perform C" and "if A and not B then perform C," which means that the hypothetical prior art does not read on them, since the hypothetical prior art explicitly teaches "if A and B then perform D," in contradistinction to the claimed invention.

The following Venn diagram summarizes the discussion made above, where the solid circle is the hypothetical prior art, the larger hollow circle represents claims 1 and 25 prior to amendment, the smaller hollow circle represents claims 1 and 25 after amendment, and the outlines of the hollow circles specify the scope of claims 1 and 25 as represented by these circles:

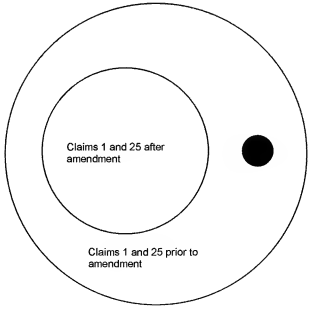

Thus, the scope of claims 1 and 25 has changed as a result of adding the phrase "even if B" to the phrase "if A then perform C." Prior to amendment, the scope of claims 1 and 25 did encompass the hypothetical prior art, whereas after amendment, the scope of claims 1 and 25 did not

encompass the hypothetical prior art. Therefore, the phrase "even if B" cannot be said to not further limit these claims, because it does.

Claim rejections under 35 USC 103(a)

All the pending claims have been rejected under 35 USC 103(a) as being unpatentable over Kiick (2003/0200250) in view of Rowlands (2003/0200250), Fischer (6,438,672), Drottar (6,170,025), Agatsuma (7,237,099), and Chi (6,209,086). Claims 1 and 25 are independent claims, from which the remaining claims ultimately depend. Applicant submits that the independent claims as previously presented are patentable, such that the remaining claims that are still pending and that have been rejected on this basis are patentable at least because they depend from independent base claims. Throughout the following discussion, claim 1 is discussed as representative of all the independent claims insofar as the rejection over Kiick in view of one or more other references in combination is concerned.

Applicant in particular now discusses two limitations of the invention that Applicant submits are not taught, disclosed, or suggested by the prior art in combination as relied upon by the Examiner.

(1) First limitation not in the prior art

Applicant notes that the claimed invention is limited to a very particular way in which interrupts for I/O devices are assigned among the nodes of a system. You have a first node to which the I/O device generating an interrupt is connected, and you have a second node at which the interrupt service routine for handling this interrupt is located. It is noted that the first node is "different" than the second node; i.e., they are not the same node. In the claimed invention, then, "if the first node . . . has a cache, memory, and at least one processor," then you assign the interrupt "to the first node." However, "if the first node does not have a cache, memory, and at least one processor, but the second node . . . does have a cache, memory, and at least one

processor," then you assign the interrupt "to the second node." Furthermore, "if both the first node and the second node do not each have a cache, memory, and at least one processor," then you assign the interrupt "to a third node . . . having memory and at least one processor." It is noted that the third node is "different" than the first and the second nodes; i.e., they are not the same nodes.

To make this ordering even more clear, Applicant amended the claimed invention in the previous office action response as follows. Thus, "if the first node . . . has a cache, memory, and at least one processor," then you assign the interrupt to the first node "*even if the second node has a cache, memory, and at least one processor*." Likewise, "if the first node does not have a cache, memory, and at least one processor, but the second node . . . does have a cache, memory, and least one processor," then you assign the interrupt to the second node "*even if a third node has a cache, memory, and at least one processor*." In this way, the first node has priority over the second node as to interrupt assignment, and the second node has priority over any third node as to interrupt assignment, where the first, second, and third nodes each has a cache, memory, and at least one processor. For instance, if all three of the first, second, and third nodes have cache, memory, and at least one processor, then in the claimed invention, you always assign the interrupt to the first node instead of to the second node and instead of to the third node. Likewise, if just the second and third nodes have cache, memory, and at least one processor, then in the claimed invention, you always assign the interrupt to the second node instead of to the third node.

It is noted that the Examiner did not actually even discuss these "even if" limitations in the obviousness rejections in the final office action, such that there is no way that the prior art in combination can be considered as rendering the claimed invention *prima facie* obviousness. "All words in a claim must be considered in judging the patentability of that claim against the prior art." (In re Wilson, 424 F.2d 1382, 1385, 165 USPQ 494, 496 (CCPA 1970)) Applicant respectfully submits that the Examiner has ignored this basic precept of obviousness analysis. It appears, rather, that the Examiner has distilled the claimed invention down to its "gist," which is

improper. "Distilling an invention down to the 'gist' or 'thrust' of an invention disregards the requirement of analyzing the subject matter 'as a whole.'" (W.L. Gore & Associates, Inc. v. Garlock, Inc., 721 F.2d 1540, 220 USPQ 303 (Fed. Cir. 1983))

In particular, in rejecting this aspect of the claimed invention, the Examiner has stated that paragraph 34 of Kiick "describes that interrupts should be assigned to the 'closest' processors," and the Examiner "interprets this to mean the interrupts for the I/O devices *should be assigned to nodes to which they are connected or to nodes where the ISRs for the said I/O device reside*." (Final office action, p. 4.) The Examiner also notes that "it would be obvious to assign the interrupt to a node which has a processor and memory, since these elements are required to service an interrupt." (Id., p. 5, relying upon Agatsuma, Drottar, and Fischer in particular.) On this basis it appears that the Examiner has concluded that the above-identified limitations of the claimed invention are taught, suggested, or disclosed by the prior art in combination.

Applicant disagrees, however, that these references in combinations rise to the level of disclosing, teaching, or suggesting the claim limitations in question. It is noted that there is a very particular order to which interrupt assignment is performed in the invention, as has been described above: if the (first) node to which the I/O device generating the interrupt has cache/memory/processors, then the interrupt is assigned to this node; if not, but if the (second) node having the interrupt service routine for this interrupt has cache/memory/processors, then the interrupt is assigned to this node; if not, then you assign the interrupt to a (third) node that does have memory/processors. Even in light of the teachings of the prior art ascribed by the Examiner, however, the prior art in combination does not teach, disclose, or suggest this particular order in assigning an interrupt.

For instance, the Examiner has stated that Kiick assigning interrupts to the closest processors means that they "should be assigned to nodes to which they are connected or to nodes where the ISRs for the said I/O device reside." However, even if this is the case, this does not mean that the node to which the I/O device is connected has priority over the node at which the

ISR resides in receiving the interrupt, assuming both have cache/memory/processors. For example, consider the following scenario: the (first) node to which the I/O device is connected has cache/memory/processors, and the (second) node at which the ISR resides also has cache/memory/processors. In the claimed invention, there is a definitive selection of which node should have the interrupt assigned to it – the first node.

By comparison, in the prior art as combined by the Examiner, there is no definitive selection of which node should have the interrupt assigned to it. Kiick is unhelpful, because as interpreted by the Examiner, Kiick merely says that the interrupt should be assigned to either the first node or the second node. Agatsuma, Drottar, and Fischer are unhelpful, because as interpreted by the Examiner, they merely say you want to select a node to assign the interrupt to that has cache/memory/processors, and in our example, both the first and the second nodes have cache/memory/processors. Thus, against these teachings, disclosures, and suggestions of the prior art, one of ordinary skill within the art does not come away with the claimed invention. In other words, the prior art does not give preference to the first node in this example, where both the first node and the second node each has cache/memory/processors.

Stated another way, the teaching of Kiick that you want to assign an interrupt to the node that is "closest," in combination with the teachings of Agatsuma, Drottar, and Fischer that you a node can only service an interrupt if it has cache/memory/processors, does not teach, disclose, or suggest the claimed invention. If both the node to which the I/O device is connected and the node having the ISR are "closest" per the interpretation of Kiick by the Examiner, and if both these nodes have cache/memory/processors such that they can service the interrupt per the interpretation of Agatsuma, Drottar, and Fischer by the Examiner, then the end result is that the prior art in combination seemingly results in the conclusion that *the prior art does not dictate to which of these two nodes should have the interrupt assigned thereto*. By comparison, the claimed invention does dictate to which of these two nodes the interrupt should be assigned: the node to

which the I/O device is connected always has the interrupt assigned to it unless this node does not have cache/memory/processors.

For this reason alone, Applicant respectfully submits that not all claim limitations are taught, suggested, and disclosed by the prior art in combination, rendering the claimed invention non-obvious and patentable over the prior art in combination.

(2) Second limitation not in the prior art

Applicant also notes that the claimed invention is limited to assigning interrupts in a round-robin manner, to dynamically modifying assignments of the interrupts based on actual performance characteristics, and for dynamically modifying assignments of interrupts that are performance critical. The Examiner has stated that the prior art in combination teaches, discloses, or suggests these limitations at least insofar as Kiick in paragraphs 25, 26, 28, and 31 teaches, discloses, or suggests these limitations. (Office action, pp. 4-5.) Applicant respectfully disagrees.

Kiick does not actually discuss assigning interrupts and dynamically modifying interrupt assignments in these excerpts, such that the prior art in combination does not teach, disclose, or suggest these limitations. Rather, Kiick discusses assigning *interrupt service routines* (ISR's), and dynamically modifying assignments of *ISR's*. (See, e.g., paragraphs 25, 26, 28, and 31, which all discuss assignment and assignment modification of ISR's, where paragraph 24 notes that "processing of each I/O interrupt is facilitated by means of an interrupt service routine (ISR) that is assignable to a specific processor.") This is a key distinction.

The claimed invention, for instance, is limited to there being a second node that has an interrupt service routine, thus distinguishing between such an ISR and an interrupt, which is described in the claim language as being assigned, or having its assignment dynamically modified. Furthermore, the prior art itself distinguishes between an ISR and an interrupt, where paragraph 24 of Kiick notes that, as has been alluded to above, "processing of each I/O interrupt is facilitated by means of an interrupt service routine (ISR) that is assignable to a specific

processor." Therefore, one cannot interpret the assignment of interrupts and the dynamic modification of interrupt assignments as in the claimed invention as being the same thing as the assignment of interrupt service routines and the dynamic modification of such interrupt service routines. Both the claimed invention and the prior art distinguish between an interrupt itself and the interrupt service routine (ISR) that services the interrupt.

As such, the limitations of the claimed invention relating to assigning interrupts and dynamically modifying interrupt assignments are not taught, disclosed, or suggested by the prior art in combination. The relied upon prior art does not actually disclose, teach or suggest assigning interrupts and dynamically modifying interrupt assignment, but rather discloses, teaches, and suggests assigning interrupt service routines and dynamically modifying interrupt service routine assignments. This is different. For at least this reason as well, then, the claimed invention is patentable over the prior art in combination.

Conclusion

Applicants have made a diligent effort to place the pending claims in condition for allowance, and request that they so be allowed. However, should there remain unresolved issues that require adverse action, it is respectfully requested that the Examiner telephone Mike Dryja, Applicant's representative, at 425-427-5094, so that such issues may be resolved as expeditiously as possible. For these reasons, and in view of the above amendments, this application is now considered to be in condition for allowance and such action is earnestly solicited.

Respectfully Submitted,

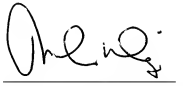

September 9, 2008
Date

Michael A. Dryja, Reg. No. 39,662
Attorney/Agent for Applicant(s)

Law Offices of Michael Dryja
1474 N Cooper Rd #105-248
Gilbert, AZ 85233
tel: 425-427-5094
fax: 425-563-2098